CLAIMS

What is claimed is:

1. A system, comprising:

an interface logic configured to pre-configure a topology of nodes to communicate via a preferred networking protocol;

a mapping logic operably connected to the interface logic, the mapping logic being configured to produce a mapping between a resource located on a first node and a port located on the first node, to selectively provide to a second node a mapping data that describes the mapping, and to selectively establish a connection that facilitates the second node accessing the resource through the port using the preferred networking protocol; and

a connection management logic operably connected to the mapping logic and the interface logic, the connection management logic being configured to control whether the mapping logic will provide the mapping data and establish the connection.

2. The system of claim 1, where to pre-configure the topology of nodes the interface logic acquires a node identifier that facilitates recording whether a node is a member of a pre-configured topology, acquires a topology configuration choice data concerning how the pre-configured topology is to be configured, pre-configures the topology based, at least in part, on the node identifier and the topology configuration choice data, and provides a topology data concerning the topology to a member of the topology.

3. The system of claim 2, where the connection management logic exerts its control based, at least in part, on the topology data and a node identifier received from the second node.

4. The system of claim 2, where a node identifier comprises one or more of, an Internet Protocol (IP) address, a value stored in one or more of a network interface card (NIC) hardware, firmware, and software, a value stored in one or more of a remote direct memory access (RDMA) NIC (RNIC) hardware, firmware, and software, a password, and a value stored on a universal serial bus (USB) token.

5. The system of claim 1, where the configuration choice data is received from one or more of, a human user via a graphical user interface (GUI), a scripting-based system, and a policy-based system.

6. The system of claim 2, where to pre-configure the topology of nodes, the interface logic determines which nodes are members of the topology, establishes a preferred computer networking protocol to be employed by members of the topology, establishes a preferred path to be employed for data communications between members of the topology, establishes a fallback networking protocol to be employed by members of the topology, and establishes a fallback path to be employed for communications between members of the topology.

7. The system of claim 2, where the topology data describes one or more of, which nodes are members of the topology, a preferred computer networking protocol to be employed between members of the topology, a preferred path to be employed for communications between members of the topology, a fallback networking protocol to be employed between members of the topology, and a fallback path to be employed for communications between members of the topology.

8. The system of claim 1, where the interface logic is further configured to control one or more resource control actions including, enabling a protocol off-load capability, aging off-loaded connections, converting idle connections to a non-off-load mode, and converting connections between an RDMA and a non-RDMA mode.

9. The system of claim 1, where the mapping logic comprises a port mapper configured to listen on a well-known port for one or more of, a request for mapping data, and a connection request.

10. The system of claim 1, the mapping logic being further configured to facilitate establishing a fallback connection between the first node and the second node according to a second networking protocol, the second networking protocol being different from the first networking protocol, where the second node may request the fallback connection after the mapping logic has been controlled to not provide the mapping data to the second node or the

mapping logic has been controlled to prevent the establishment of a connection between the first node and the second node using the first networking protocol.

11. The system of claim 10, the connection management logic being configured to block access to a first resource on the first node via the preferred networking protocol and to permit access to a second resource on the first node via a fallback networking protocol.

12. The system of claim 1, where the resource supports one or more of, remote direct memory access (RDMA) between the first node and the second node, and protocol off-loading at the first node.

13. The system of claim 1, where one or more of, the interface logic, the mapping logic, and the connection management logic are located on one or more of, a network interface card (NIC), and a remote direct memory access (RDMA) NIC (RNIC).

14. The system of claim 2, where the connection management logic exerts its control based on analyzing the topology data and one or more of, time of day, network traffic, load, and resource availability.

15. The system of claim 1, where the connection management logic operates at a session layer associated with the first networking protocol.

16. The system of claim 15, where the first networking protocol includes a Transmission Control Protocol (TCP) transport layer and an Internet Protocol (IP) network layer.

17. A computer configured with a pre-configured topology connection management system, the system comprising:

an interface logic configured to pre-configure a topology of nodes to communicate via a preferred networking protocol or a fallback networking protocol, where to pre-configure the topology of nodes the interface logic acquires a node identifier that facilitates recording whether a node is a member of a pre-configured topology, acquires a topology configuration choice data concerning how the pre-configured topology is to be configured, pre-configures the topology based, at least in part, on the node identifier and the topology configuration

choice data, and provides a topology data concerning the topology to a member of the topology;

a mapping logic operably connected to the interface logic, the mapping logic being configured to produce a mapping between a resource located on a first node and a port located on the first node, to selectively provide to a second node a mapping data that describes the mapping between the resource and the port, and to selectively establish a connection between the first node and the second node, where the connection facilitates the second node accessing the resource through the port using the preferred networking protocol; and

a connection management logic operably connected to the mapping logic and the interface logic, the connection management logic being configured to control whether the mapping logic will provide the mapping data to the second node, and whether the mapping logic will establish the connection, where the connection management logic exerts its control based, at least in part, on the topology data and a node identifier received from the second node.

18. A method, comprising:

acquiring a set of node identifiers associated with nodes to be considered for inclusion in a pre-configured topology of nodes that can communicate within the topology using a preferred computer networking protocol;

establishing the pre-configured topology of nodes; and

making available a membership data concerning the pre-configured topology of nodes.

19. The method of claim 18, where the set of node identifiers are acquired from one or more of, a human user through a graphical user interface (GUI), a scripting-based system, and a policy-based system.

20. The method of claim 19, where a node identifier comprises one or more of, an Internet Protocol (IP) address, a value stored in a network interface card (NIC) hardware, a value stored in a NIC firmware, a value stored in a NIC software, a value stored in a remote direct memory access (RDMA) NIC (RNIC) hardware, a value stored in an RNIC firmware, a

value stored in an RNIC software, a password, and a value stored on a USB (Universal Serial Bus) token.

21. The method of claim 18, where establishing the pre-configured topology of nodes includes:

determining node membership in the pre-configured topology;

establishing a preferred computer networking protocol to be employed by members of the topology;

establishing a preferred path to be employed for communications between members of the topology;

establishing a fallback computer networking protocol to be employed by members of the topology;

establishing a fallback path to be employed for communications between members of the topology; and

recording the topology membership, preferred computer networking protocol, preferred path, fallback computer networking protocol, and fallback path in the membership data.

22. A computer-readable medium storing processor executable instructions operable to perform a method, the method comprising:

acquiring a set of node identifiers associated with nodes to be considered for inclusion in a pre-configured topology of nodes that can communicate within the topology using a preferred computer networking protocol or a fallback computer networking protocol;

establishing the pre-configured topology of nodes, where establishing the pre-configured topology of nodes includes determining node membership in the pre-configured topology, establishing a preferred computer networking protocol to be employed by members of the topology, establishing a preferred path to be employed for communications between members of the topology, establishing a fallback computer networking protocol to be employed by members of the topology, establishing a fallback path to be employed for communications between members of the topology, and recording the topology membership, preferred computer networking protocol, preferred path, fallback computer networking protocol, and fallback path in the membership data; and

making available a membership data concerning the pre-configured topology of nodes.

23. A method, comprising:

acquiring a set of node identifiers associated with nodes to be considered for inclusion in a pre-configured topology of nodes that can communicate within the topology using a preferred computer networking protocol;

establishing the pre-configured topology of nodes;

distributing a membership data concerning the pre-configured topology of nodes to nodes that are in the pre-configured topology of nodes;

selectively adding or deleting a node from the pre-configured topology of nodes and, in response to selectively adding or deleting the node, redistributing the membership data; and

selectively managing a computer networking resource, and in response to selectively managing the computer networking resource, redistributing the membership data.

24. The method of claim 23, where a node identifier comprises one or more of, an Internet Protocol (IP) address, a value stored in a network interface card (NIC) hardware, a value stored in a NIC firmware, a value stored in a NIC software, a value stored in a remote direct memory access (RDMA) NIC (RNIC) hardware, a value stored in an RNIC firmware, a value stored in an RNIC software, a password, and a value stored on a USB token.

25. The method of claim 23, where establishing the pre-configured topology of nodes includes:

determining node membership in the pre-configured topology;

establishing a preferred computer networking protocol to be employed by members of the topology;

establishing a preferred path to be employed for communications between members of the topology;

establishing a fallback computer networking protocol to be employed by members of the topology;

establishing a fallback path to be employed for communications between members of the topology; and

recording the group membership, preferred computer networking protocol, preferred path, fallback computer networking protocol, and fallback path in the membership data.

26. The method of claim 23, where selectively managing a computer networking resource includes one or more of, enabling a protocol off-load capability, aging an off-loaded connection, converting an idle connection to a non-off-load mode, and converting a connection between an RDMA mode and a non-RDMA mode.

27. A method, comprising:

in a first node, receiving from a second node, via an open computer networking protocol, a request to establish a connection between the first node and the second node via the open computer networking protocol, where the connection facilitates the second node accessing a resource associated with the first node;

determining whether the second node is a member of a pre-configured topology that includes the first node by examining a node identifier associated with the second node; and

selectively not establishing the connection between the first node and the second node via the open computer networking protocol if it is determined that the second node is not a member of the pre-configured topology that includes the first node.

28. The method of claim 27, where the method is performed by a session layer logic associated with the open computer networking protocol.

29. The method of claim 27, where the open computer networking protocol includes a Transmission Control Protocol (TCP) transport layer and an Internet Protocol (IP) network layer.

30. A computer-readable medium storing processor executable instructions operable to perform a method, the method comprising:

in a session layer logic in a first node, receiving from a second node, via an open computer networking protocol that includes a Transmission Control Protocol (TCP) transport layer and an Internet Protocol (IP) network layer, a request to establish a connection between the first node and the second node via the open computer networking protocol, where the connection facilitates the second node accessing a resource associated with the first node;

determining whether the second node is a member of a pre-configured topology that includes the first node; and

selectively not establishing the connection between the first node and the second node via the open computer networking protocol if it is determined that the second node is not a member of the pre-configured topology that includes the first node.

31. A method, comprising:

in a first node, receiving from a second node a mapping request for a mapping data that describes a relationship between a resource on the first node and a port on the first node;

selectively providing the mapping data to the second node based on determining that the second node is a member of a pre-configured topology that includes the first node by examining a node identifier associated with the second node;

receiving from the second node a connection request to establish a connection between the first node and the second node via a first networking protocol, where the connection facilitates accessing the resource;

selectively establishing the connection based on determining that the second node is a member of a pre-configured topology that includes the first node by examining a node identifier associated with the second node; and

via a second networking protocol, receiving from the second node a fallback connection request to establish a fallback connection between the first node and the second node, where the fallback connection request requests that the fallback connection be established via the second networking protocol, where the fallback connection granted in response to the third request will not provide access to the resource via the first networking protocol.

32. The method of claim 31, where selectively establishing the connection is based additionally on an availability of the resource.

33. The method of claim 31, where the resource is located on one or more of, a network interface card (NIC), and a remote direct memory access (RDMA) NIC (RNIC) associated with the first node.

34. The method of claim 31, where the resource supports one or more of remote direct memory access (RDMA) between the first node and the second node, and protocol off-loading at the first node.

35. The method of claim 31, where the first networking protocol includes a Transmission Control Protocol (TCP) transport layer and an Internet Protocol (IP) network layer.

36. A computer-readable medium storing processor executable instructions operable to perform a method, the method comprising:

in a first node, receiving from a second node a mapping request for a mapping data that describes a relationship between a resource on the first node and a port on the first node;

selectively providing the mapping data to the second node based on determining, by examining a node identifier associated with the second node, that the second node is a member of a pre-configured topology that includes the first node;

receiving from the second node a connection request to establish a connection between the first node and the second node via a first networking protocol, where the connection facilitates accessing the resource;

selectively establishing the connection based on determining that the second node is a member of a pre-configured topology that includes the first node by examining a node identifier associated with the second node; and

via a second networking protocol, receiving from the second node a fallback connection request to establish a fallback connection between the first node and the second node, where the fallback connection request requests that the fallback connection be established via the second networking protocol, where the fallback connection granted in response to the third request will not provide access to the resource via the first networking protocol.

37. A system, comprising:

means for determining whether a client node is a member of a pre-configured topology to which a server node belongs;

means for rejecting a request that will lead to the undesired consumption of a server resource if the requesting client node is not a member of the pre-configured topology to which the server node belongs; and

means for establishing a connection between the client node and the server node using a networking protocol preferred by members of the pre-configured topology.

38. A set of application programming interfaces embodied on a computer-readable medium for execution by a computer component in conjunction with pre-configured topology connection management, comprising:

a first interface for communicating a group data configured to facilitate determining whether a client node is a member of a pre-configured topology to which a server node belongs; and

a second interface for communicating a resource management data that facilitates determining whether a client node will be granted a connection to a resource located on a topology member node.